IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

Date: May 19, 2003
Express Mail: ER216392591US

In re application of: **J. P. Brezin, et al**

Serial No.: **09/282,860**

Filed: **March 31, 1999**

Docket No.: **YO999-121**

Board of Patent Appeals and Interferences
Washington, D.C. 20231

TRANSMITTAL OF APPEAL BRIEF UNDER 37 CFR 1.192

Sir:

Transmitted herewith, in triplicate is an Appeal Brief with respect to the Notice of Appeal filed March 18, 2003 for the above-identified patent application.

This Appeal Brief is being filed on behalf of other than a small entity.

Authorization is given to charge amount of $320.00, for filing a Brief in support of appeal in accordance with 37 CFR 1.17(f), to Deposit Account 50-0510. A duplicate copy of this authorization is enclosed.

The Assistant Commissioner is hereby authorized to charge any required additional fee, and charge back any overpayment, to Deposit Account No. 50-0510.

Respectfully submitted,
J. P. Brezin, et al

Anne Vachon Dougherty
Registration No. 30,374
Tel. (914) 962-5910

#15

RECEIVED
2003 MAY 20 PM 1:20
BOARD OF PATENT APPEALS
AND INTERFERENCES

Serial No. 09/282,860
Art Unit No. 2172

I HEREBY CERTIFY THAT THIS CORRESPONDENCE IS
BEING DEPOSITED IN EXPRESS MAIL AS
EXPRESS MAIL: ER216392591US
ON DATE OF DEPOSIT: May 19, 2003
PERSON DEPOSITING: ANNE VACHON DOUGHERTY

Anne Vachon Dougherty 5/19/03
Signature & Date

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

In Re Application of	:	May 19, 2003
J. P. Brezin, et al	:	Group Art No.: 2172
Serial No. 09/282,860	:	Examiner: J. Fluerantin
Filed: March 31, 1999	:	for IBM Corporation Anne Vachon Dougherty
Title: OPTIMIZATION OF SYSTEM PERFORMANCE BASED ON COMMUNICATION RELATIONSHIP		3173 Cedar Road Yorktown Heights, NY 10598

RECEIVED
MAY 21 2003
Technology Center 2100

Board of Patent Appeals and Interferences
Washington, D.C. 20231

APPEAL BRIEF (37 CFR 1.192)

Applicant hereby appeals to the Board of Patent Appeals and Interferences from the decision dated November 18, 2002 of the Primary Examiner finally rejecting Claims 1-6, 11, 23-24, 26-31, 37-41 in the above application. Applicant

09282860
05/22/2003 SCHAPMAN 00000003 500510
01 FC:1402 320.00 CH

respectfully request that the Board of Patent Appeals and Interferences consider the arguments presented herein and reverse the Examiner's rejection.

I. REAL PARTY IN INTEREST

The appeal is made on behalf of Applicants who are real parties in interest with respect to the subject patent application.

II. RELATED APPEALS AND INTERFERENCES

There are no pending related appeals or interferences with respect to the subject patent application.

III. STATUS OF CLAIMS

There are forty-one (41) claims pending in the subject patent application, numbered 1-41. Claims 1-6, 11, 23-24,

26-31, 37-41 stand rejected. Claims 7-10, 12-22, 25, 32, 34-36 are objected to as being dependent upon a rejected base claims, and the Examiner has indicated that they would be allowable if rewritten in independent form including all of the limitations of the base claim and any intervening claims. Claim 33 was not enumerated as either rejected or objected to in the Final Office Action.

IV. STATUS OF AMENDMENTS

The status of the prosecution of the application is as follows:

December 3, 1999	-	Preliminary Amendment filed.
July 5, 2001	-	Office Action rejecting all claims.
November 5, 2001	-	Amendment filed introducing claim amendments to the independent claims.
February 12, 2002	-	Final Office Action finally rejecting all claims
May 13, 2002	-	CPA filed

June 3, 2002	-	Office Action rejecting Claims 1-6, 11, 23-24, 26-31, 33, and 37-41
September 3, 2002	-	Amendment filed introducing claim amendments to the independent claims.
November 18, 2002	-	Final Office Action, finally rejecting Claims 1-6, 11, 23-24, 26-31, and 37-41.
March 18, 2003	-	Notice of Appeal filed

V. SUMMARY OF INVENTION

The subject invention is a method and a program storage device for implementing a method for improving performance in a user's communication system by automatically extracting and integrating relationship information and constructing a relationship data structure for use in automatically modifying a user's information retrieval query based on the relationship data structure. The relationship information is extracted from multiple heterogeneous information sources, including such sources as previous communications by the user, previous communications to the user, user

calendar entries, address book, and relationship databases, such as organization charts. The extracted relationship information is then integrated into and stored in a relationship data structure. When a user inputs an information retrieval query (e.g., to send an e-mail meeting notice to a particular person), the relationship data structure is used to automatically modify the query (e.g., to add additional addressees to the e-mail meeting notice).

VI. STATEMENT OF ISSUES OF APPEAL

There following issues are on appeal:

(1) whether the Examiner erred in omitting Claim 33 from either the rejected claims or the objected-to claims; and

(2) whether the Examiner erred in holding that the teachings of the Paul patent render the claimed invention obvious.

VII. GROUPING OF CLAIMS

The rejected Claims can be considered in the following groups for purposes of this appeal:

(I) Group I: Claims 1-5, 33, 37-39, and 40-41 reciting the method of extracting and integrating relationship information, building and storing a relationship data structure, and query modification;

(II) Group II: Claims 6 and 11, further reciting the use of preferences in the handling of relationship information;

(III) Group III; Claims 23, 26, further reciting resolving ambiguities in relationship information;

(IV) Group IV; Claim 24, further reciting recommending communications channels based on relationship information;

(V) Group V; Claims 27-29, further reciting detecting and handling changes to relationship information; and

(VI) Group VI; Claims 30-31, further reciting the organization of the relationship data structure.

VIII. ARGUMENT

ARGUMENT (1)

With regard to issue (1), whether the Examiner erred in omitting Claim 33 from either the list of rejected claims or the list of objected-to claims; Applicants submit the following remarks.

To address a typographical error in the original patent application, which resulted in two different claims bearing number "33", Applicants submitted a Preliminary Amendment canceling the first of the two claims numbered "33" and adding Claim 41 to recite the subject matter of the first of the two claims numbered 33. In three subsequent Office Actions, the Examiner rejected Claim 33. However, in the most recent Final Office Action from which Applicants appeal, the Examiner omitted Claim 33 from consideration without explanation. Applicants believe that the Examiner overlooked Claim 33 when reviewing the case and inadvertently considered it to be a canceled claim. Applicants respectfully assert that the Final Office Action should not stand due to this error.

ARGUMENT (2)

As to issue (2), whether the Examiner erred in holding that the teachings of the Paul patent render the claimed invention obvious; Applicants submit the following arguments.

The subject application teaches and claims a method to optimize information retrieval based on communication relationships which, as recited in the two independent claims, comprises the steps of automatically extracting and integrating relationship information from multiple heterogeneous information sources; automatically building and storing a relationship data structure to represent the relationship information; and automatically modifying an information retrieval query based on the relationship data structure (Claims 1 and 40). All of the remaining claims depend either directly or indirectly from Claim 1 and recite further limitations thereto. Specifically, the Group II Claims 6 and 11, further recite the use of preferences in the handling of relationship information; the Group III Claims 23, 26, further recite resolving ambiguities in relationship information; the Group IV Claim 24, further

recites recommending communications channels based on relationship information; the Group V Claims 27-29, further recite detecting and handling changes to relationship information; and, the Group VI Claims 30-31, further recite the organization of the relationship data structure.

The Paul patent is directed to an electronic mail (hereinafter "e-mail") filtering system which compares content (e.g., the "SUBJECT:" or "FROM:" entries) of an incoming e-mail message to a list of terms on an exclusion list. If the incoming e-mail content matches any term on the exclusion list, the e-mail is marked as "junk mail". The Paul exclusion list can be updated by information from external sources, such as the so-called spam probes placed strategically in the network. The spam probes are homogeneous "sources" which convey address information to be added to the exclusion list. In addition, the Paul patent system can generate a spam alert signal to warn other user terminals or network servers of detection of spam e-mail.

In distinguishing the presently claimed invention from the Paul system, Applicants first point to the first claim step of extracting and integrating relationship information from multiple heterogeneous information sources. Applicants

refer the Examiner to the Specification (e.g., page 7, line 12-page 8, line 12) where the term "relationship information" is detailed. As a first argument, Applicants assert that the Paul patent does not teach or suggest "relationship information from multiple heterogeneous information sources". The cited teachings of the Paul patent, at Col. 6, lines 17-21, do not teach or suggest multiple heterogeneous sources, do no teach or suggest communications relationship information, and do not teach or suggest steps of extracting and integrating communications relationship information from multiple heterogeneous information sources.

The present invention extracts communications relationship information, as detailed on pages 7 and 8, from multiple heterogeneous sources (including calendars, network information, organization charts, e-mail databases, etc.) in order to facilitate information retrieval queries (and resulting communications) using that relationship information. The present invention does not seek to exclude information/communications like Paul, but rather seeks to facilitate information retrieval and communications.

The primary Paul teachings relate to examining the "FROM" field of an incoming e-mail or the text of the e-mail and comparing it to entries in an exclusion list. The "FROM" or text information is not extracted from the e-mail, nor is that information integrated with any other information from any other sources. The text is simply compared. Clearly, therefore, the Paul step of comparing would not suggest extracting and integrating as claimed. Further, the Paul step of flagging or marking the e-mail as "junk" would not suggest extraction and integration of relationship information from the e-mail. The Paul step does not integrate information from multiple heterogeneous information sources...it simply flags one e-mail as "junk".

The Examiner has argued on page 4 of the Office Action that the Paul patent generating of a spam alert is "readable as optimize information retrieval based on communication relationship" and is "readable as automatically extracting and integrating relationship information from multiple heterogeneous information sources." Applicants respectfully disagree. The Paul patent does not make any mention of information retrieval in either the filtering of incoming mail or in the generating of an alert signal. Moreover, the

updating of an exclusion list with information from a received spam alert is not the same as nor suggestive of automatically extracting and integrating relationship information from multiple heterogeneous sources. The Paul system simply receives a spam alert from one of a plurality of homogeneous sources (i.e., an e-mail from a trusted location) and adds the listed spam information to its exclusion list. Applicants believe that the Paul teachings regarding updating an exclusion list with information from homogeneous sources is not the same as nor suggestive of automatically extracting and integrating relationship information from multiple heterogeneous information sources. Relationship information is not extracted by the user site from a source under the Paul teachings. Rather, exact e-mail address/source information is delivered to the user site. Further, the e-mail address information which is provided to the user site is not relationship information, it is address information. Again Applicants note that "relationship information" was expressly defined for the present invention and does not simply mean "anything related to a communication". Finally, Applicants again assert that the Paul updating of an exclusion list with an e-mail

address from a homogeneous source is not the same as or suggestive of integrating relationship information from multiple heterogeneous sources.

With regard to the second claim feature of automatically building and storing a relationship data structure to represent the relationship information, the Examiner states on page 5 that "Paul does not explicitly indicate..." that feature. The Examiner goes on to conclude, however, that "Paul indicates the spam control center automatically analyzes the received mail to identify the source of the message extracts and processes the source data from the received message, and generates an alert signal containing the processed source data the alert signal may also contain filtering instructions used to enable network servers and user terminals to automatically detect spam this alert signal is broadcast to all network servers of all user terminals or both within the communications network a filtering system implemented at the servers or user terminals automatically receives the alert signal automatically updates stored filtering data using the source data retrieved from the alert signal and automatically controls delivery of subsequently-received e-mail messages

from the identified spam source; which is readable as automatically building and storing a relationship data structure to represent the relationship information *(sic)*." Insofar as Applicants can understand the Examiner's points, given the poor sentence structure and lack of punctuation, Applicants respectfully disagree. Applicants respectfully assert that the Examiner is using the same Paul teachings to negate two distinct steps of the invention (extracting and integrating in one step and building and storing in the next distinct step), when clearly Paul is only performing one function. Further, the alert signals are used for updating a list with the new e-mail/source address information from one of multiple homogeneous spam probe sources. Again Applicants argue that Paul does not have multiple heterogeneous sources; does not obtain relationship information from its sources; does not extract any information from the sources, but simply receives an alert with address information and/or filtering information from those sources; and, does not build a relationship data structure, but simply adds an entry to a list. Clearly the Paul patent is neither teaching nor suggesting building a

relationship data structure when it teaches that a list is maintained.

The Examiner further notes that the Paul exclusion list processor may store an exclusion list for each e-mail address or alternatively an exclusion list for each group of e-mail addresses organized by domain or other group. Applicants disagree with the Examiner's conclusion that organizing an exclusion list in such a fashion is the same as or suggestive of building a relationship data structure representing relationship information extracted and integrated from multiple heterogeneous information sources.

The Examiner acknowledges that Paul does not teach the building of a relationship data structure and yet concludes that it would be obvious to modify Paul to include that step. Applicants respectfully argue that, absent some teachings which would motivate one skilled in the art to so modify the Paul patent, the Examiner cannot conclude that such would be an obvious "leap". The Paul patent makes no mention of relationship information, let alone of using relationship information, or of building a relationship data structure.

Finally, Applicants argue that there is nothing in Paul which teaches or suggests the step of modifying an information retrieval query based on the relationship data structure. Paul is not directed to information retrieval and makes no mention of query processing in any context. Clearly, therefore, one would not logically modify Paul to include queries and the modification of queries based on a relationship data structure. The Examiner has attempted to equate the Paul filtering of incoming e-mails using an exclusion list as query processing and to equate the updating of the exclusion list as modifying a query. Applicants have amended the independent claims to highlight the fact that the query which is being modified using relationship information is a query which the user generates for information retrieval. What is claimed is not a "query" to filter incoming user mail, it is an information retrieval query. The claim step is not directed to modifying an exclusion list, but to modifying a query based on information in a created relationship data structure to expedite processing of the query.

The Examiner states that, in the Paul patent, "a user terminal filtering application 200 for use in the present

invention (*sic, the Paul* system) includes an exclusion list manager 202 for creating, storing and automatically maintaining a user exclusion list the user exclusion lists preferably includes all identification data needed to determine the status of incoming e-mail messages data in the exclusion list may be divided into categories corresponding to the fields of incoming e-mail messages; which is readable as which is readable as automatically modifying an information retrieval query based on the relationship data structure" (*sic*). Applicants again note that the Examiner's sentence structure and lack of punctuation make it difficult to discern the points being made. Applicants do understand that the Examiner is here trying to equate the Paul teachings of the system looking at fields of an incoming e-mail and comparing information in those fields to the stored entries on an exclusion list to the claimed feature of automatically modifying an information retrieval query. Applicants respectfully point out that Paul is not performing any information retrieval query processing. The Paul patent does not make any mention of queries or of information retrieval. Applicants believe that the Examiner

has erred in trying to extend the Paul teachings to include information retrieval query processing.

Finally, the Examiner concludes that it would have been obvious for Paul to build and store a relationship data structure. Applicants respectfully disagree. What the Paul patent teaches, in the cited passages at Columns 8 and 9, is the e-mail filtering using an exclusion list, or a plurality of exclusion lists, one for each e-mail address if implemented at a server. There is simply nothing in the Paul patent which teaches or suggests that relationships between prior and present communications be detected by extraction and integration of relationship information. Further, there is no suggestion that a relationship data structure be created based on such extraction and integration. Finally, there is nothing to suggest that queries be modified at all, let alone based on relationship information in a relationship data structure which has been built and stored based on the extraction and integration of information from prior communications.

Applicants respectfully assert that the terms which are used in the pending claims are not arbitrarily chosen, nor can they be arbitrarily re-interpreted by the Examiner. The

claimed relationship information is defined and explained in the Specification (see: e.g., page 7, line 12-page 8, line 12; and page 10, line 8-page 20, line 7) and cannot simply be re-defined by the Examiner to mean merely "entries on a manually-created user exclusion list". Paul's exclusion list is NOT a relationship data structure and the Paul entries are not relationship information. Accordingly, one skilled in the art would not leap to the conclusion that Paul's system for comparing incoming text to a list of entries could be modified to automatically extract relationship information from prior communications and proceed to integrate, build, store, and use such relationship information, let alone modify future communications/queries based on that relationship information. Applicants respectfully request that the Board consider the language of the independent claims, Claims 1 and 40, in light of the Specification, and overturn the rejections.

With respect to the other rejected claims, Applicants respectfully assert that the Paul patent neither teaches nor suggests the invention as claimed. As noted above, all of the rejected dependent claims, Claims 2-6, 11, 23-24, 26-31,

37-39 and 41, depend from independent Claim 1. To facilitate understanding of the issues, Applicants have grouped the claims as set forth above. Applicants believe that all of the claims are patentable over the Paul patent teachings.

With respect to the remaining Group I claims, Claims 2-4, 37-39 and 41, Applicants note that those claims all relate to more details regarding modifying the query. Applicants respectfully repeat the assertion that the Paul patent does not teach querying, let along modifying a query by prioritizing and filtering the retrieval of information (Claim 2), the augmenting of information from the sources (Claim 3), modifying to optimize delivery of query results (Claim 4), creating sub-queries (Claim 37), aggregating results from sub-queries (Claim 38), excluding results from sub-queries (Claim 39), and prioritizing and filtering a list of address mappings to facilitate sending e-mails (Claims 33 and 41). The Examiner has analogized the Paul adding of information to exclusion lists to the Claim 2 language of prioritizing and filtering retrieval of information. Applicants respectfully disagree. Prioritizing relationship information and filtering

relationship information retrieved for an information retrieval query is clearly not suggested by updating an exclusion list of addresses.

Further, the augmenting information from multiple heterogeneous sources is clearly not suggested by updating exclusion lists. What is claimed by Claim 3 is augmenting relationship information which has been received. The invention as claimed adds to the extracted information, it does not just add entries to a list. With regard to Claim 4, the Paul system for excluding e-mails clearly does not teach or suggest modifying a query to optimize delivery of query results. As noted above, the Paul patent seeks to halt delivery, not to optimize it. Further, there is no query in the Paul patent for which delivery of results can be optimized. With respect to Claim 5, the claim recites the multiple heterogeneous sources, which clearly are not taught or suggested by Paul. As to the sub-query language of Claims 37-39, the cited teachings at Col. 8, lines 37-41 regarding multiple exclusion lists clearly do not apply since no query is being processed, let alone a sub-query. Finally, with regard to Claim 41 (and presumably Claim 33), the cited Paul passage from Col. 2, lines 17-24 teaches

discarding junk e-mail and clearly does not teach prioritizing and filtering a list of name-to-mail address mappings to facilitate sending e-mail.

With regard to the Group II claims, Claims 6 and 11 further recite the use of preferences in the handling of relationship information. Applicants again respectfully note that the Paul patent does not teach relationship information, let alone the assigning different preferences to the heterogeneous information sources (Claim 6) and the assigning of weight to each information source based on a preference; and computing the aggregate communication intensity, based on the weight and the preference (Claim 11). While the Paul patent takes into account user preferences when building the exclusion lists (e.g., "I prefer not to accept e-mails that offer mortgages"), that is not the same as or suggestive of assigning different preferences to multiple heterogeneous sources (e.g., "relationship information extracted from previous meeting notices takes precedence over relationship information extracted from old organization charts") for Claim 6 and the user thereof for Claim 11. The passages cited against Claims 6 and 11 describe comparing incoming e-mails to the

exclusion lists, wherein the exclusion lists have been created with user preference data, which is clearly not the same as or suggestive of the claimed invention.

With regard to the Group III, Claims 23 and 26, which further recites resolving ambiguities in relationship information, Applicants reiterate that the Paul patent does not teach relationship information, let alone a step of resolving name ambiguity by using the relationship from heterogeneous information sources to determine one or more of an e-mail address, phone number, and a full name (Claim 23) or detecting inconsistency among data in the heterogeneous information sources (Claim 26). The Examiner cites a Paul passage at Col. 4, lines 22-34 against Claim 23, however, that cited passage just describes a spam probe, which clearly does not obviate resolving name ambiguity with relationship information from multiple heterogeneous sources. Similarly, the passage cited against Claim 26, at Col. 2, lines 25-49 teaches details of the spam probe and spam alerts but does not teach or suggest detecting inconsistency among data from multiple heterogeneous sources.

With regard to the Group IV claim, Claim 24, further recites recommending communications channels based on relationship information. The Paul patent does not teach or suggest collecting or using relationship information. Furthermore, recommending communications channels for communications that are predicated on a modified information retrieval query is clearly not taught or suggested by Paul. The cited passage from Col. 5, lines 10-20 of the Paul patent details examining an incoming e-mail to compare fields in the e-mail to the exclusion list. There is no discussion of recommending a communication channel for communications based on relationship information.

With regard to the Group V Claims 27-29, which further recite detecting and handling changes to relationship information, Applicants respectfully assert that the cited passage, again from Col. 2, lines 25-49, teaches details of the spam probe and spam alerts but does not teach or suggest detecting changes in relationship information (Claim 27), propagating the changes (Claim 28) or alerting the changes (Claim 29).

With regard to Group VI, Claims 30-31, which further recite the organization of the relationship data structure,

Applicants have maintained that the Paul patent simply does not teach the building or maintaining of a relationship data structure. Clearly the Examiner could not point to teachings which then provide details of how to organize the data structure. Rather, the Examiner has cited passages which describe a network for implementing the Paul e-mail monitoring. A network having nodes (i.e., user terminals, etc.) is clearly not what is being claimed by the present recitation of integrating relationship information using a relationship data structure graph with links (Claim 30) and intensity vectors (Claim 31).

CONCLUSION

Applicants respectfully assert that the Examiner has erred in omitting Claim 33, and in rejecting Claims 1-6, 11, 23-24, 26-31, and 37-41 as unpatentable over the Paul patent. In light of the foregoing arguments, Applicants request that the decision of the Examiner, rejecting Claims 1-6, 11, 23-24, 26-31, 37-41, be overturned by the Board and that those claims, along with Claim 33 and the allowable yet objected to Claims 7-10, 12-22, 25, 32, 34-36, be passed to issuance.

Respectfully submitted,
J. Brezin, et al

By: 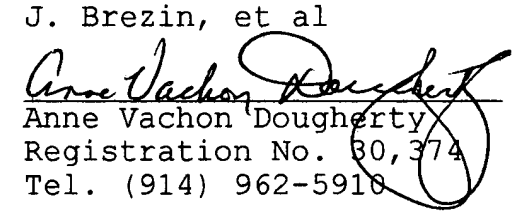
Anne Vachon Dougherty
Registration No. 30,374
Tel. (914) 962-5910

APPENDIX OF CLAIMS

CLAIMS UNDER APPEAL IN BOLD TYPE

ALLOWABLE CLAIMS IN NORMAL TYPE

1. A method to optimize information retrieval based on communication relationships, comprising the steps of:

automatically extracting and integrating relationship information from multiple heterogeneous information sources;

automatically building and storing a relationship data structure to represent the relationship information; and

automatically modifying an information retrieval query based on the relationship data structure.

2. A method of claim 1, wherein said step of modifying a query comprises the steps of prioritizing and filtering the retrieval of related information.

3. A method of claim 1, wherein said step of modifying a query comprises the steps of augmenting information from the heterogeneous information sources.

4. A method of claim 1, wherein said step of modifying a query comprises the step of modifying a query to optimize delivery of query results.

5. A method of claim 1, wherein the heterogeneous information sources are selected from the group consisting of one or more of: people-managed data sources; organization charts; mailing lists; calendar entries; personal address books; priority lists of contacts; and automated system log type information including phone logs and e-mail logs.

6. A method of claim 1, further comprising the step of assigning different preferences to the heterogeneous information sources.

7. A method of claim 1, further comprising the steps of:

said step of building a data structure further comprising the step of tracking communication intensities

between each pair of communication entities via each information source; and

integrating the relationship information from the heterogeneous information sources, in response to said tracking step.

8. A method of claim 7, further comprising the step of:

deriving a relation-group for each communication entity based on a pre-specified criterion on said communication intensities.

9. A method of claim 8, further comprising the step of:

selecting relation-group entities of a communication entity based on aggregate communication intensities to the communication entity.

10. A method of claim 9, further comprising the step of: computing an aggregate communication intensity from an entity A to an entity B based on a weighted sum of the communication intensities from said entity A to said entity B via each information source.

11. A method of claim 6, further comprising the steps of:

assigning a weight to each information source based on a preference; and

computing the aggregate communication intensity, based on the weight and the preference.

12. A method of claim 8, further comprising the step of:

deriving relation-group entities of an entity allowing one or more of an indirect relationship and an inferred relationship.

13. A method of claim 12, further comprising the step of:

said deriving step further comprising the step of deriving a relation-group of an entity *A*, which can include the relation-group entities of an entity in the relation-group of the entity *A*.

14. A method of claim 8, further comprising the step of driving an awareness service based on a relation-group relationship.

15. A method of claim 7, further comprising the step of:

building and maintaining additional persistent data structures based on the results of the query to facilitate the response on future queries, based on the relationship data structure.

16. A method of claim 15 wherein the additional persistent data structure can be a personal address/phone book based on the communication intensity.

17. A method of claim 7, further comprising the step of determining a significance of a relationship between two entities.

18. A method of claim 17, further comprising the step of determining the significance of a relationship based on the aggregate communication intensity.

19. A method of claim 7, wherein the tracking step can be subject based.

20. A method of claim 8, wherein the relation-group can be subject based.

21. A method of claim 18, wherein the significance of the relationship can be subject based.

22. The method of claim 18, further comprising the step of downloading information based on the significance of the relationship.

23. The method of claim 3, further comprising the step of resolving name ambiguity by using the relationship from the heterogeneous information sources to determine one or more of an e-mail address, phone number, and a full name.

24. The method of claim 4, further comprising the step of recommending a communication channel based on a recipient characteristic.

25. The method of claim 7, further comprising the step of caching a document and information based on the significance of the relationship.

26. The method of claim 1, further comprising the step of detecting inconsistency among data in the heterogeneous information sources.

27. The method of claim 26, further comprising the step of detecting changes in the relationship information maintained.

28. The method of claim 27, further comprising the step of propagating the changes.

29. The method of claim 27, further comprising the step of alerting the changes.

30. A method of claim 1, further comprising the steps of:

integrating the relationship information from the multiple heterogeneous sources using a graph wherein each

node represents a communication entity, and a link between a pair of nodes represents the existence of a communication relationship between the two nodes.

31. A method of claim 30, further comprising the step of labeling each link with a communication intensity vector, where each dimension of the communication intensity vector represents a communication intensity from an information source.

32. A method of claim 12, further comprising the step of calculating aggregate communication intensities taking into account the indirect relationship.

33. A method of claim 3, further comprising the step of obtaining relevant information from the heterogeneous information sources, said information selected from the group consisting of one or more of: phone numbers; e-mail addresses: mailing addresses; office location; department; or manager, from various information sources.

34. A method of claim 7, further comprising the step of calculating a communication intensity based on a number of communication events.

35. A method of claim 7, further comprising the step of calculating a communication intensity based on both a number of communication events and their temporal characteristics.

36. A method of claim 7, further comprising the step of calculating a communication intensity based on an analysis of a content of a communication event.

37. A method of claim 3, further comprising the step of modifying the query to create one or more sub-queries.

38. A method of claim 37, further comprising the step of aggregating results from the sub-queries.

39. A method of claim 37, further comprising the step of excluding results from the sub-queries.

40. A program storage device readable by a machine, tangibly embodying a program of instructions executable by the machine to perform method steps for optimize information retrieval based on communication relationships, said method steps comprising:

automatically extracting and integrating relationship information from multiple heterogeneous information sources;

automatically building and storing a data structure to represent the relationship information; and

automatically modifying an information retrieval query based on the relationship data structure.

41. A method of claim 2, further comprising the step of prioritizing and filtering a list of name-to-e-mail address mappings to facilitate sending e-mail.